

The Undeﬂated Domain of Semantics
Paul M. Pietroski, University of Maryland

It is, I suppose, a truism that an adequate theory of meaning for a natural language L will associate each sentence of L with its meaning. But the converse does not hold. A theory that associates each sentence with its meaning is *not*, by virtue of that fact, an adequate theory of meaning. For it is also a truism that a semantic theory should explain the (interesting and explicable) semantic facts. And one cannot decree that the relevant facts are all reportable with instances of schemata like ‘S means that p’ or ‘S, by virtue of its meaning, is true iff p’. Investigation suggests that there is *much* more for semanticists to explain: natural languages exhibit synonymies, ambiguities, and entailments; for any string of words, there are endlessly many meanings it cannot have; there are semantic generalizations, including crosslinguistic generalizations, that go uncaptured and unexplained by merely associating sentences with their meanings; *etc.* Initially, one might think these facts are “peripheral” and can thus be ignored if the aim is to explain why sentences mean what they do. But the study of natural language suggests otherwise. (One can’t tell, in advance of investigation, which facts are peripheral to a given domain. It was initially tempting to think that one could ignore falling bodies, and the tides, if the aim was to explain why planets move as they do.)

We *find out* what a theory of meaning can and should explain by *doing* semantics. We start by trying to explain a range of facts pretheoretically regarded as semantic; and then we see where inquiry leads. This would be too obvious to mention, much less write about, except that some philosophers seem to have lost track of the point. And this is not harmless ignorance of empirical research. For it has been coupled with the observation that one can, without deploying the theoretical apparatus that semanticists standardly deploy, specify trivial algorithms that associate each sentence of a language with its meaning. Thus, confusion about the aims of semantics—and failure to consider its successes—is easily transformed into skepticism about the need for nontrivial theories of the sort that semanticists try to offer. So I will be stressing that facts like those reported with

(1) ‘Dogs bark’ means that dogs bark

reflect the tip of an iceberg.

Another truism is that speakers of English understand English sentences. If you know English, then modulo performance limitations, you can associate each sentence of English with its meaning. But it doesn’t follow that if you can associate each sentence of English with its meaning, you thereby understand English (or its sentences)—at least not if understanding is taken to be the natural phenomenon exhibited by native speakers. Speakers of a language recognize various synonymies, ambiguities, and entailments; they know what certain strings of words cannot mean; they acquire semantic knowledge at a characteristic (and quick) pace across linguistic environments; *etc.* Investigation suggests that these capacities are not peripheral. When a speaker of English grasps the meaning of ‘Dogs bark’, she does so by exercising a competence that lets her do more than simply discern the meanings of English sentences. Describing the nature of this competence, and how speakers exercise it, is a matter for empirical investigation. So one cannot decree that understanding a sentence S is simply a matter of figuring out what S means—or figuring this out by some stipulated method. One can define ‘understanding’ this way; but then it remains a hypothesis, to be evaluated in the light of *all* relevant evidence, that anyone who understands a sentence comprehends it in the way that native speakers do. Chomsky (1977, 1986, 2000) and others—e.g., Higginbotham (1985), Larson and Segal (1995)—have made these points repeatedly. But let’s try again.

1. Horwich (1997, 1998) advocates the following thesis—page references are to Horwich (1997):

Understanding one of one’s own complex expressions (nonidiomatically) is, by definition, nothing over and above understanding its parts and knowing how they are combined (p. 504).

If ‘one has worked out how a certain sentence is constructed from primitive syntactic elements’, then ‘provided one knows the meanings of those elements’ one understands the sentence

automatically and without further ado...No further work is required; no further process needs to be involved, leading from these initial conditions to the state of understanding the sentence (p. 504).

Correlatively, the fact reported with (1)—or using Horwich’s convention of ‘capitalizing an English expression to obtain a name of its meaning’, the fact that ‘dogs bark’ means DOGS BARK—is ‘constituted by whatever is the fact regarding its mode of construction and the meanings of its constituents’. According to Horwich, this is the fact that the sentence ‘results from putting words meaning what “dog” and “bark” mean into a schema meaning what “*ns v*” means’ (p. 505).

Horwich uses ‘*ns v*’ to refer to the sentential schema ‘*noun(+plural) verb*’; and while I’m not entirely sure what ‘what “*ns v*” means’ means, the idea is that the meaning of any sentence that instantiates the schema consists in (i) its having the syntax in question and (ii) its words, which instantiate parts of the schema in a certain order, having the meanings they do. Thus, he writes:

...the meaning property

x means DOGS BARK

consists in what I shall call the “construction property”

x results from putting terms whose meanings are DOG and BARK, in that order, into a schema whose meaning is *NS V* (p. 505).

I have quoted at length because this is an audaciously simple—and so audaciously strong—thesis about the meanings of natural language sentences (and what it is to understand a sentence). Those familiar with the literature in semantics might have doubted that anyone would advocate such a stark view.

1.1 Horwich is to be commended for offering such a clear statement of a simple view which, if correct, would be very attractive. But it would have been better, had he offered some evidence for the proposal. Nowhere does Horwich try to show how the claims above would help explain the *range* of facts that semanticists regularly discuss. The focus is almost exclusively on the mere fact that sentential meanings are compositional.¹ And even here, matters are far more complex than he suggests.

In a footnote (citing Jim Higginbotham), Horwich considers the possibility that the syntax of a sentence *S* may itself make a substantive contribution to the meaning of *S*. If this is so, then understanding *S* would require *more* than ‘simply understanding the words in it and seeing how they have been combined’. But according to Horwich,

any so-called “method of combination” that intuitively has a meaning (for example, predication or conjunction or ‘*ns v*’) can be regarded as a schematic *constituent* of the sentence. Thus, we can make it a matter of stipulation that no meaning attaches to the procedures by which these and other constituents may be combined (n2, p. 505).

I have three objections to this reply. First, I don’t see how the *syntax* of a sentence can be regarded as a *constituent* of a sentence. Consider

(2) John ran slowly

whose syntax is roughly as follows: [_S [_{NP} John][_{VP} [_V ran][_{Adv} slowly]]]. If someone says that combining ‘ran’ with ‘slowly’ signifies a certain semantic operation, over and above the meanings of the words combined, I don’t see how one can gloss this suggestion by saying that the relevant aspect of natural language syntax—[_{VP} [_V ...][_{Adv} ...]]—is a constituent of (2).

Second, the available evidence suggests that Horwich is wrong. A large and growing body of work suggests that the meaning of (2) should be represented along the lines of

(2M) $\exists e[\text{Agent}(e, \text{John}) \ \& \ \text{Ran}(e) \ \& \ \text{Slow}(e)]$

where ‘Ran’ and ‘Slow’ are predicates of events and ‘Agent’ expresses a *thematic* relation that holds between an event done by someone (or something) and the relevant person (or thing).² If this is correct, then a natural thought is that the conjunctive aspects of (2M), indicated by the ampersands, are reflected by the syntactic structure of (2). On this view, natural language concatenation is a way of expressing *predicate-conjunction*; see Pietroski (forthcoming-a, b) for defense and further discussion. But perhaps

natural language syntax corresponds instead to *function-application*, as in familiar views deriving from the work of Frege (1891) and Montague (1970). This works nicely for sentences like

(3) John ran

involving just predicates and arguments.³ Though sentences involving adjuncts, like (2) and

(4) The doctor from Chicago is from Chicago

present a challenge. The two tokens of ‘from Chicago’ in (4) cannot express functions of the same type: the first token combines with ‘doctor’ to form a complex unary predicate; the second token *is* a unary predicate. Thus, some sort of *type-shifting* is required; and it is hardly obvious that (if this semantic approach is correct) speakers appreciate the semantic effect of type-shifting, in sentences with adjuncts, simply by appreciating the syntax of such sentences. Indeed, a standard thought is that type-shifting reflects agrammatical aspects of meaning due to the intrusion of extralinguistic cognitive factors into (what is fundamentally) a function-argument semantics. One might, however, adopt a mixed view: some but not all syntax corresponds to function-application; see Heim and Kratzer (1998). In any case, questions about the semantic contribution of syntax are subtle and hard.

Third, and most importantly, this is not a matter for *stipulation*. One has to *find out* what natural language syntax contributes to the meanings of natural language expressions. One may as well stipulate that no meaning attaches to auxiliary verbs, or words beginning with ‘q’. As Horwich notes, one could modify his thesis by saying that ‘the meaning of a sentence consists in its mode of constitution having a certain meaning and its constituents having certain meanings’ (p. 505). But this raises the question, which lies at the heart of compositional semantics, of *which* meanings various aspects of syntax have. (And correlatively, which *types* of meanings words have.) Horwich seems to think he can sidestep such questions by saying that a given mode of combination means what it does: [_{VP} [_V ...][_{adv} ...]] means [_{VP} [_V ...][_{ADV} ...]]. But if we just homophonically report the meanings of constituents and syntax, then unsurprisingly, we fail to explain the facts that motivate adoption of more substantive proposals.

For example, if (2) is true, then so is (3). Speakers of English know that the inference from (2) to (3) is impeccable. Many eventish semanticists try to explain this fact, and a *host* of related facts, by saying that (2) and (3) involve covert quantification over events as in (2M) and

(3M) $\exists e[\text{Agent}(e, \text{John}) \ \& \ \text{Ran}(e)]$.

The inference from (2M) to (3M) is impeccable because it is an instance of conjunction reduction in the scope of an existential quantifier; and the hypothesis is that speakers who understand (2) and (3) tacitly grasp this. Perhaps this is the wrong explanation; maybe a type-shifting proposal, or something else entirely, accounts for the relevant facts. But all Horwich can say is that (2) means what it does because of what its parts mean and how those parts are arranged, and similarly for (3). This does not even begin to explain why these sentences have meanings such that the truth of the former guarantees the truth of the latter—or how speakers of English know this. The syntax of [_{VP}[_V ran][_{Adv} slowly]] does not reveal that any individual who satisfies this verb-phrase *also* satisfies [_V ran].

This point is not confined to eventish inferences involving adjuncts. Horwich-style theories provide no explanation for the *many* entailment patterns that semanticists regularly discuss. If every dog barked, and Fido is a dog, then Fido barked. But the fact that ‘every dog barked’ means what it does, because of what its constituents mean and how they are arranged, does *not* help us explain why (speakers know that) this sentence is true iff $\{x:\text{dog}(x)\} \subseteq \{x:\text{barked}(x)\}$. Without some such (nonhomophonic and substantive) representation of what the quantified sentence means, we cannot capture what seems to be a striking relation between meaning and implication. Saying that ‘every’ means EVERY just doesn’t do it. We need a theory that spells out *why* the quantified sentence is true iff a subset relation holds between the set of things that satisfy ‘dog’ and the set of things that satisfy ‘bark’. One could go on for weeks, as one does in an undergraduate semantics course, about the ubiquity of this point. If Pat boiled the soup, the soup boiled. If John thinks that Bob likes himself, John thinks that Bob likes *Bob*; cf. ‘John thinks that Bob likes him’. If every kid swam, every tall kid swam; whereas if most kids swam, it doesn’t follow

that most tall kids swam. If every kid is a kid who swam, every kid swam. Why? And how do speakers know such things?⁴

So far as I can tell, Horwich is committed to denying that these facts are due to the meanings of the natural language sentences; although since he never mentions such facts, it is a little difficult to know what his view is. One can introduce a technical term ‘meaning*’ and stipulate that certain facts are not due to the meanings* of sentences. But then the question is whether the notion of meaning* is of any interest. Either way, work is required: one has to show that the facts are not, *pace* standard theories, best explained as reflections of meanings; or one has to show that the notion of meaning* is the theoretical notion we should deploy in this domain.

1.2 Semanticists also try to account for certain *contrasts* in meaning; and a Horwich-style semantics will not help us account for the interesting facts that have been discovered. Consider the following sentences, variants of which are discussed by Higginbotham (1983) and Vlach (1983):

- (5) I heard Pat sing
- (6) I heard Pat sang.

In (5), where the embedded verb is untensed, ‘Pat’ occupies a semantically transparent position; if Pat is Chris, and (5) is true, then I heard Chris sing. But (6) is truth-conditionally equivalent to

- (7) I heard that Pat sang

in which ‘Pat’ occupies a semantically opaque position. Why? Obviously, (5) and (6) differ syntactically: the former has an untensed embedded clause, while the latter has a tensed embedded clause (and probably a covert complementizer). But why is this difference correlated with such a significant semantic difference? It does no good to say that the arrangement of words in (5) means what it means, and similarly for (6). We want to know why these sentences have different *kinds* of meanings; and the first step is to provide *illuminating* representations of what they do mean. One wonders if Horwich thinks the opacity of (7) can be explained by saying that ‘that Pat sang’ means THAT PAT SANG. For if more explanation than this is required, I don’t see how his proposal provides it—or even allows for it.

Or consider cases of ambiguity. The string of words in

- (8) You cannot stop a philosopher with a theory
- can support (at least) two different meanings, much like
- (9) I saw the philosopher with binoculars.

The two meanings of (8) are correlated with different syntactic structures, as indicated below

- (8a) <stop{a [philosopher (with a theory)]}>
- (8b) {[stop(a philosopher)](with a theory)}.

But why do these different arrangements of the words in (8) have *different* meanings? Why is one reading roughly synonymous with ‘stop a philosopher *who has* a theory’, while the other is roughly synonymous with ‘stop a philosopher *by using* a theory’. Merely saying that (8a) means 8A, while (8b) means 8B doesn’t tell us why 8A ≠ 8B. But this is a crucial part of what is to be explained. The force of the point increases exponentially with a multiply ambiguous string of words like

- (10) I can duck and hide whenever visiting relatives might scare me.

Cases of *nonambiguity* illustrate the flip side of the same point. Speakers of English cannot use

- (11) Was the child who lost kept crying

to ask whether the child who *was* lost kept crying. Instead, (11) can only be used to ask whether (perversely) the child who lost *was* kept crying. The standard explanation is that the syntax of (11) involves a covert trace of ‘was’ between ‘lost’ and ‘kept’; and since natural languages do not permit extraction of auxiliary verbs out of a relative clause, (11) cannot be the result of moving ‘was’ from a position between ‘who’ and ‘lost’.⁵ But this raises a question. Why *doesn’t* the structure indicated by

- (11a) Was_i [_S [_{NP} the child who lost]_j [_{VP} t_i kept crying]]

support the impermissible meaning? Why does the relevant sentential frame have the meaning it does, as opposed to other possible meanings—like the perfectly coherent meaning that (11) cannot exhibit? If our

semantic theory has the consequence that the VP in (11a) is satisfied by an individual x iff x was kept crying, and that this severely constrains what (11a) can mean, we have at least the beginning of an answer; but not so, if we simply say that the sentential frame means what it does. Similar questions arise with respect to many of Chomsky's famous examples. Consider

- (12) John ate an apple (14) John is too clever to catch a fish
 (13) John ate (15) John is too clever to catch

Why *can't* (15) mean that John is too clever to catch something or other? Why does it have to mean, on *disanalogy* with (13), that John is too clever for us to catch him? Saying that (14) and (15) mean what they do doesn't even describe—much less explain—the facts of interest here.⁶

Continuing with the golden oldies, the contrast exhibited by

- (16) John is eager to please
 (17) John is easy to please

is especially striking. While (16) can only mean that John is eager that *he* please relevant parties, (17) can only mean that it is easy for relevant parties to please John. This suggests the following syntactic structures, with italicized items and coindexing representing covert elements and cointerpretation:

- (18) John_{*i*} is [eager [*e_i* to please *e_j*]]
 (19) John_{*j*} is [easy [*e_i* to please *e_i*]]

Presumably, 'eager' and 'easy' differ semantically in ways that explain (in conjunction with various syntactic facts) why 'John' cannot be the object of the embedded sentence in (18), while 'John' cannot be the subject of the embedded sentence in (19). And presumably, 'ready' differs yet again, since

- (20) John is ready to eat

is ambiguous. But what are the relevant facts about the meanings of 'eager' and 'easy' (and 'ready')? It does no good to say that 'eager' means EAGER and 'easy' means EASY. If these are trivial claims—'eager' and 'easy' mean what they do—then Horwich's theses about meaning (and understanding) imply that the facts just illustrated are *not* reflections of the meanings of sentences like (18) and (19). In which case, an alternative explanation of the facts is owed. On the other hand, if it takes work to specify the meanings of 'eager' and 'easy' in ways that help explain the (nondisquotational) semantic facts, then a "theory" that reports these meanings by simply capitalizing the verbs is inadequate.

1.3 One could keep making such points *ad nauseum*. But let me conclude this section with one last remark. As the examples above suggest, a great deal of work on natural language semantics concerns the role of covert elements. Horwich does not address this issue. But one wonders what his view is. Are covert elements to be understood as sentential *constituents*? Should we say that *e* means *E*—or perhaps that '*e_i*' means *E₁*? Or are covert elements to be understood as aspects of sentential frames? Do all covert elements contribute to meaning in the same way? What *do* the embedded clauses in (18) and (19) mean? Are these live issues, to be decided by actually studying natural languages, or are they matters for stipulation? Given the voluminous body of work on anaphora, quantification, binding, and various sorts of empty categories, it is hard to believe that anyone would advance Horwich's thesis about what it is to understand a sentence without discussing these matters.

2. Which leads one to wonder if Horwich really wants to advance a thesis about understanding. As I mentioned in note 1, his main concern is to show that one can explain the compositionality of meaning in the (self-avowedly) trivial manner he suggests. For if this is correct, the fact that sentential meanings are compositional does not impose substantive constraints on the meanings of words (or syntax). And if semantic theories do not impose such constraints, Horwich's (1990) "deflationary" account of truth is not threatened by the prospect of a semantic theory that reveals substantive constraints on our notion of truth.

2.1 It's not obvious, however, that Horwich's proposal even explains what it is designed to explain. The fact reported by

- (1) 'Dogs bark' means that dogs bark

may follow from the facts he cites: the plural noun 'dogs' means DOGS; the verb 'bark' means BARK;

and the sentential frame ‘plural-noun verb’ means what it does. But not every case of deduction is an explanation; subsuming an instance under a generalization doesn’t always explain the instance. One might worry that Horwich’s schemata *summarize* their instances, as opposed to answering why-questions (and in that sense rendering their instances more comprehensible). Showing that a generalization G is *explanatory* usually takes the form of showing that G figures in a wide *range* of explanations, and that G interacts with other generalizations to explain various *complex* phenomena. Horwich does none of this. But set such concerns aside, and grant that there is a thin but legitimate sense in which his proposal explains correspondingly thin facts about what English sentences mean.

And let us not concern ourselves with indexicality, which threatens to render the scope of disquotational schemata vanishingly small. And let us not dwell on the fact that if our metalanguage is English, then we cannot just capitalize to report semantic facts like the following: ‘Masaaki-ga migite-o ageta’ means, in Japanese, that Masaaki raised his right arm. And let’s not ask why the case markers ‘ga’ and ‘o’ are semantically inert, in a way that causative morphemes (in Japanese and other languages are not). And let’s not worry about how to report—much less explain—the meaning of the Edo sentence ‘Òzó ghà lè èvbàré khién’, which transliterates as ‘Ozo will cook food sell’. It seems to mean *something* like: Ozo will perform a complex action, unified by a single plan, that starts with a cooking of some food and ends with a selling of that very food.⁷ But let’s suppose that simply deriving the Edo equivalent of

(21) ‘Ozo cook food sell’ means OZO COOK FOOD SELL

explains why ‘Òzó ghà lè èvbàré khién’ means what it does.

One might wonder, though, if there is more to semantics than explaining why a sentence means what it does. Perhaps one can report the meaning of a sentence in some nondisquotational but structured way, and then explain why the sentence has that meaning (so described) as opposed to others, with the aim of uncovering the *richness* of understanding. But set all this aside and grant that a Horwich-style theory can explain *the* sense in which meaning is compositional. Or put another way, grant that facts about mere compositionality do not themselves provide reason to reject a Horwich-style theory in favor of a more substantive semantic theory.

Horwich makes it clear, in several places, that this is his central point. For example, when he considers modifying his stated thesis to allow for the possibility that syntax itself contributes to meaning, he notes that ‘the explanation of compositionality would be no less trivial’ (n.2, p. 505). And he speaks of his general strategy as ‘deflationary’ because

...it shows that the compositionality of meaning is much easier to explain than we have often been led to believe. It would not seem to be the case, as contended by Davidson and his many followers, that compositionality dictates an explication of meaning properties in terms of *reference* and *truth conditions* (p. 507).

Elsewhere, he cites the ‘Davidsonian’ thesis that ‘compositionality can be explained *only* by explicating meanings as truth conditions’ (p. 519, my italics). But Horwich provides no citations for this view, which he thinks is ‘widely accepted’ (p. 517); and in my opinion, no sane Davidsonian would accept it.

Many (neo-)Davidsonians hold—on the basis of *all* the available evidence—that speakers are able to understand sentences because speakers tacitly know various truth-theoretic axioms from which they can extract consequences concerning the truth-conditions of sentences.⁸ If the only argument for this view were that meaning (and/or understanding) is compositional, then Horwich would be right to say that a very strong thesis had been based on very slender evidence. But this strikes me as a bizarre interpretation of the Davidsonian program, especially as developed in recent years; see Evans (1981), Higginbotham (1985, 1986), Taylor (1985), Parsons (1990), Schein (1993), Larson and Segal (1995), *etc.* A plethora of facts, including the elementary examples reviewed in section one, suggest that natural language predicates have substantive satisfaction conditions.

By contrast, Horwich’s own proposal is targeted exclusively to compositionality. So even if he is right that the mere compositionality of meaning does not motivate adoption of a nondeflationary

semantic theory, the interest of this point is purely hypothetical. If the totality of available evidence suggests that a deflationary theory is woefully inadequate to the semantic facts as we find them, it hardly matters if such a theory would have been adequate had there been fewer facts. If the phenomena semanticists explain can be explained without resort to nondeflationary theoretical apparatus, that is of interest. But we need to see the more parsimonious explanations, in the context of specific examples of the phenomena in question.

Otherwise, the deflationary proposal threatens to degenerate into the following truism: if (i) compositionality is the only phenomenon that a theory of meaning has to explain, and (ii) compositionality is a “thin” phenomenon that can be explained by a deflationary theory of meaning, then (iii) there is no reason for preferring a nondeflationary semantic theory to a deflationary one. This conditional is interesting only in the presence of reasons for thinking that (i) and (ii) are true. But (ii) is dubious at best; and (i) is false if the actual study of natural language semantics is any indication. One can claim that an entire research program is misguided, and that the facts semanticists study do not bear on the nature of meaning. But this is the sort of thing that needs showing, example by example. Alternatively, one can invent technical terms like ‘meaning*’ (and ‘understanding*’) and make the following stipulations: (i*) compositionality is the only phenomenon that a theory of meaning* has to explain, and (ii*) compositionality is a “thin” phenomenon that can be explained by a deflationary theory of meaning*. Then by stipulation, (iii*) there is no reason for preferring a nondeflationary semantic* theory to a deflationary one. But it is hard to see why this tautology bears on the study of meaning.⁹

In fact, if one is going to impose restrictions on which facts get to constrain theories of meaning, it seems arbitrary to let compositionality (and not, say, nonambiguity) count as relevant. A simpler claim would be that the schema

(22) ‘p’ means that p,

or some variant that accommodates indexicality, already constitutes an adequate theory of meaning. So far as I can tell, the only merit of Horwich’s proposal over this one is that Horwich offers a (thin) explanation for *one* sort of semantic fact. But I can’t see why (mere) compositionality should be *the* semantic fact of interest. One can, of course, develop whatever theory—subject to whatever constraints—one chooses. But one doesn’t get to stipulate that

Understanding one of one’s own complex expressions (nonidiomatically) is, by definition, nothing over and above understanding its parts and knowing how they are combined (p. 504). And a deflationary account of compositionality hardly shows that the nondisquotational semantic facts don’t impose substantive constraints on the meanings of words and syntax.

3. Finally, having raised the issue, let me briefly address the related suggestion that natural languages don’t have a compositional semantics—except perhaps in some trivializing sense; see Schiffer (1992, 1994a, 1994b, 1998).¹⁰ If the *only* argument for (substantive) semantic compositionality in natural languages was that speakers can understand arbitrarily many sentences, one might suggest the following alternative explanation: there are endlessly many sentences; and a speaker understands a sentence ‘p’ of her language if she appreciates the truth of some suitable (noncompositional) schema like (22). But there are *many* arguments for semantic compositionality in natural language; some have been reviewed here.

Consider, once again,

- (8) You cannot stop a philosopher with a theory
- (11) Was the child who lost kept crying
- (11a) Was_i [_S [_{NP} the child who lost]_i [_{VP} t_i kept crying]]
- (18) John_i is eager [_{e_i} to please e_j]
- (19) John_j is easy [_{e_i} to please e_j].

How does one explain the fact that the string of words in (8) can support the two different meanings it can support, without supposing that those meanings are determined compositionally? On standard views, the meaning of {[stop(a philosopher)](with a theory)} is determined by combining the meanings of ‘stop’

and ‘a philosopher’ *and then combining the result* with the meaning of ‘with a theory’; and similarly, *mutatis mutandis*, for the other reading of (8)—<stop{a [philosopher (with a theory)]}>. Only this latter expression has a constituent satisfied by philosophers with (i.e., philosophers who have) a theory. If this isn’t correct, then how *is* the ambiguity of (8) explained?

Prima facie, (11) cannot have the impermissible meaning—the child who was lost kept crying?—*in part because* the meaning of (11a) is partly determined by the meaning of [_{VP} t_i kept crying], which is satisfied by things kept crying. If one denies that the meaning of (11a) is determined compositionally, and in part by the quite specific meaning of the verb phrase, one owes an alternative explanation for why (11) has (only) the quite specific meaning it does. Similarly, if the meanings of (18) and (19) are not determined in part by the meaning of ‘e_i to please e_j’, how does one even *begin* to explain the semantic contrast between (18) and (19)?

If the meanings of

(2) John ran quickly

(3) John ran

aren’t compositional, why *does* the truth of (2) guarantee the truth of (3)? And more importantly, if speakers don’t understand (2) and (3) compositionally, how do speakers know that the inference from (2) to (3) is impeccable? And so on.

Indeed, it is hard to find an interesting explanation of a (nonlexical) semantic fact that doesn’t presuppose compositionality. So “the” argument that natural languages have a substantive compositional semantics is that this assumption figures in so many explanations of particular (nondisquotational) semantic facts. Correlatively, the reason for wanting more than a deflationary theory of meaning is that such a theory fails to explain just about all the interesting semantic facts that we have discovered.

References

- Baker, M. 1997: Thematic Roles and Grammatical Categories. In L. Haegeman, ed., *Elements of Grammar* (Dordrecht: Kluwer).
- Baker, M. and Stewart, O. forthcoming: On Taking the “Construction” out of Serial Verb Constructions.
- Castañeda, H. 1967: Comments. In Rescher 1967.
- Chomsky, N. 1977: *Essays on Form and Interpretation*. New York: North Holland.
- 1986: *Knowledge of Language*. New York: Praeger.
- 2000: *New Horizons in the Study of Language and Mind*. New York: Cambridge.
- Crane, S. and Pietroski, P. 2000: Nature, Nurture, and Universal Grammar.
To appear in *Linguistics and Philosophy*.
- Davidson, D. 1967: The Logical Form of Action Sentences. In Rescher 1967.
- 1984: *Inquiries into Truth and Interpretation*. Oxford: OUP.
- Evans, G. 1981: Semantic Theory and Tacit Knowledge. In S. Holzman and C. Leich, eds., *Wittgenstein: To Follow a Rule* (London: Routledge and Kegan Paul).
- Frege, G. 1891: Function and Concept. In P. Geach & M. Black, trans., *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Blackwell, 1980).
- Heim, I. & Kratzer, A. 1998: *Semantics in Generative Grammar*. Oxford: Blackwell.
- Higginbotham, J. 1983: The Logical form of Perceptual Reports. *Journal of Philosophy* 80:100-27.
- 1985: On Semantics. *Linguistic Inquiry* 16: 547-93.
- 1986: Linguistic Theory and Davidson’s Program. In E. Lepore, ed., *Truth and Interpretation*. Oxford: Blackwell.
- Horwich, P. 1990: *Truth*. Oxford: Blackwell.
- 1997: The Composition of Meanings. *Philosophical Review* 106: 503-32.
- 1998: *Meaning*. Oxford: OUP.

- Larson, R. and Segal, G. 1995: *Knowledge of Meaning*. Cambridge, MA: MIT.
- Lepore, E. and Fodor, J. 1999: Why Compositionality Won't Go Away: Reflections on Horwich's "Deflationary" Theory. Rutgers Technical Report Series, volume 46.
- Lepore, E. and Ludwig K. forthcoming: *Donald Davidson: Meaning, Truth, Language, and Reality*.
- Montague, R. 1970: English as a Formal Language. Reprinted in *Formal Philosophy*. (New Haven: Yale University Press, 1974).
- Neale, S. 1990: *Descriptions*. Cambridge, MA: MIT Press.
- Parsons, T. 1990: *Events in the Semantics of English*. Cambridge, MA: MIT.
- Pietroski, P. forthcoming a: Function and Concatenation. In G. Preyer, ed., *Logical Form* (Oxford: OUP).
 —forthcoming b: *Events and Semantic Architecture*. Oxford: OUP.
 —forthcoming c: The Character of Natural Language Semantics. In A. Barber, ed., *Epistemology of Language* (Oxford: OUP).
- Rescher, N. (ed.) 1967: *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.
- Schein, B. 1993: *Events and Plurals*. Cambridge, MA: MIT.
- Schiffer, S. 1987. *Remnants of Meaning*. Cambridge, MA: MIT.
 —1992: Belief Ascription. *Journal of Philosophy* 89: 499-521.
 —1994a: The Language-of-Thought Relation and its Implications. *Philosophical Studies* 76: 263-285.
 —1994b: A Paradox of Meaning. *Nous* 28: 279-324.
 —1998: Meanings and Concepts. *Lingue e Style* 33
 —2000: Critical Study: Horwich on Meaning. *Philosophical Studies* 50: 527-536.
- Stewart, O. forthcoming: The Serial Verb Construction Parameter. New York: Garland.
- Taylor, B. 1985: *Modes of Occurrence*. Oxford: Blackwell.
- Vlach, F. 1983: On Situation Semantics for Perception. *Synthese* 54:129-52.

Notes

1. For reasons I return to, in section two, Horwich is especially concerned to show that one can account for the compositionality of meaning without appeal to a truth-theoretic semantics of the sort urged by Davidson (1984). Horwich (1998) embeds the discussion of compositionality in a broader discussion of issues concerning meaning that philosophers often discuss; but he still fails to engage with the work done in actually constructing semantic theories. I turn to some examples presently.
2. See Castañeda's (1967) comment on Davidson (1967); see also Higginbotham (1983, 1985), Taylor (1985), Parsons (1990), Schein (1993), Larson and Segal (1995), *etc.*
3. If one associates 'ran' with a function (say from individuals to truth-values) and 'John' with an entity in the domain of that function, while associating (3) with the value of the relevant function given the relevant entity as argument, one is effectively associating concatenation with function-application. Thus, one might replace (2M) and (3M) with '[Slowly(Ran)]John' and 'Ran(John)'. But even this "minimal" contribution by syntax does not correspond to a *constituent* of (3).
4. Horwich (1998) discusses analyticity without considering such examples. Note that even if 'x boiled y' contains the word 'cause' covertly, which seems unlikely, the crucial question remains: why does [_S[_{NP} ...][_{VP} (cause) [_{XP} ...]]] imply [_{XP} ...]? The last inference pattern holds for all natural language determiners; but we can easily invent determiners for which it doesn't. While 'only' is not a determiner, it illustrates the point: only boys are boys who swam; but it hardly follows that only boys swam. See Larson and Segal (1995) for discussion.
5. Very young children are sensitive to this constraint. For discussion, and a review of other psycholinguistic data that bear on semantic theories, see Crane and Pietroski (2000).
6. Or to take an apparently mundane—but if current research is any indication, suprisingly deep—question: why does 'Brutus stabbed Caesar' have no reading on which Brutus is the stabbee and Caesar is the stabber? See the references in note 2; see also Baker (1997).
7. See Stewart (forthcoming), Baker and Stewart (forthcoming) for discussion of the data and syntax; see Pietroski (forthcoming b) for discussion of the semantics.
8. On p. 520, Horwich seems to offer an argument to the effect that this account *cannot* be right. But it relies on a rejectable premise concerning tacit knowledge. And there is no attempt to provide alternative explanations for the phenomena neo-Davidsonians regularly cite. (It would have helped, for example, if Horwich had explicitly showed how to deal with quantification and indexicality on his view; cf. Neale (1990), Larson and Segal (1995), Lepore and Ludwig (forthcoming). Davidsonians may sometimes use 'compositionality' in referring to the full range of phenomena that a compositional semantics should explain; see also Lepore and Fodor (1999) on "reverse" compositionality.
9. That said, one might wonder if semantic theories need to be theories of *truth*, in order to explain the relevant facts; see Chomsky (2000), Pietroski (forthcoming c). But this requires attention to specific examples; and I know of no alternative that is deflationary, in the way that Horwich urges, about the meaning properties of words and syntax. For example, if verbs (and adverbs) are unary predicates of events, this already goes far beyond saying that 'ran' means RAN, and that 'pushed' means PUSHED. And the eventish thesis has considerable ramifications for how other terms (and syntax) contribute to

meaning; see Pietroski (forthcoming-a,b).

10. For example, if one simply identifies the meaning of a word with its contribution (whatever it is) to the meanings of sentences in which the word appears, then the meaning of a sentence is presumably determined (given its syntax) by the meanings of its words; but this does not specify the meanings of words in any substantive way that reveals *why* sentences mean what they do. One can add that speakers map public sentences onto sentences of an internal language of thought. But this conjecture hardly explains the semantic facts, absent a detailed proposal (concerning the alleged mapping) that accounts for the kinds of phenomena noted above (and below).